arXiv:1403.5843v1 [cs.PL] 24 Mar 2014

Efficient Synthesis of Network Updates

Jedidiah McClurg[1], Nate Foster[2], and Pavol Černý[1]

[1] University of Colorado Boulder
[2] Cornell University

Abstract. Configuration updates are a leading cause of instability in networks today. A key factor that makes updates difficult to implement is that networks are distributed systems with hundreds or thousands of nodes all interacting with each other. Even if the initial and final configurations are correct, naively updating individual nodes can easily cause the network to exhibit incorrect behaviors such as forwarding loops, black holes, and security vulnerabilities. This paper presents a new approach to the network update problem: we automatically generate updates that are guaranteed to preserve important correctness properties. Our system is based on counterexample-guided search and incorporates heuristics that quickly identify correct updates in many situations. We exploit the fact that the search algorithm asks a series of related model checking questions, and develop an efficient incremental LTL model checking algorithm. We describe a prototype implementation of our system in OCaml, and present experimental results showing that it efficiently generates updates for real-world topologies with thousands of nodes.

1 Introduction

Most networks must be frequently updated, for reasons ranging from planned maintenance, to balancing traffic loads, to patching security vulnerabilities. Unfortunately, even when updates are planned in advance, they are difficult to implement correctly—even if the initial and final configurations are correct, the update process can step through incorrect intermediate configurations. Indeed, disruptions caused by updates occur often in practice, and lead to a host of problems including lost connections, decreased performance, and transient outages.

A key factor that makes updates difficult to implement stems from the fact that networks are distributed systems, with hundreds or even thousands of interacting devices. To transition a network from one configuration to another, an operator cannot simply update all devices simultaneously. Moreover, the network must continue to function while the update propagates through the system. Today, operators typically construct the low-level instructions that modify the configurations of individual devices by hand—a tedious and error-prone process.

To address these problems, the research community has developed a number of mechanisms for implementing network updates while preserving important invariants [10,15,27,32,29]. For example, *consensus routing* uses distributed snapshots to avoid routing loops during BGP reconfiguration [15], while *consistent*

updates uses versioning to ensure that every packet is processed according to the initial or final configuration [29]. But these mechanisms do not address the general problem—they are limited to specific protocols, or they come with heavy costs, requiring substantial additional space on switches or time to complete.

This paper presents a new approach to the network update problem. Rather than forcing operators to construct updates manually or use ad hoc, expensive mechanisms, we develop techniques for generating network updates automatically. Instead, an operator simply specifies the new configuration and an LTL property, and our system generates a sequence of low-level instructions that implement the update and ensure the property, or fails if no such sequence exists.

This approach enables operators to specify just the properties required for applications, which enables strategies that are ruled out in other systems (e.g., consistent updates preserve *all* safety properties [29]). If the operator specifies a trivial property, then it may be safe to update the switches in any order, but if they specify a more complex policy, such as "always traverse the firewall," then the update must be carefully constructed to avoid violations during the update.

Our system works by searching through the space of all possible sequences of individual switch updates, learning from counter-examples and employing an incremental model checking algorithm to re-use previous results whenever possible. The incremental model checking algorithm is novel in its own right—we exploit the fact that correct configurations are loop-free and obtain an efficient procedure for re-checking after changes. Because our tool poses a series of closely related model checking questions, incremental techniques yield enormous performance improvements on large examples. Furthermore, we propose and evaluate a heuristic that uses partial model checking results to avoid backtracking.

We have built a prototype implementation in OCaml, and used it to generate updates for a number of topologies, configurations, and properties including the real-world Topology Zoo dataset, and synthetic datacenter networks. Our experimental results demonstrate the effectiveness of our approach: it is able to scale to networks with thousands of nodes, and it dramatically outperforms a tool built using a state-of-the-art symbolic model checker.

To sum up, the main contributions of this paper are as follows:

- We investigate using synthesis to generate network updates (§2).
- We design a counter-example-guided search algorithm, prove it correct, and identify heuristic optimizations (§4).
- We develop an incremental LTL model checker for loop-free models (§5).
- We describe our prototype implementation in OCaml and present experimental results on realistic topologies, configurations, and properties (§6).

Overall, our system takes a challenging programming problem—building correct updates—and automates it, yielding a powerful network management tool.

2 Overview

Abstractly, a network can be thought of as a graph with switches as nodes and links as edges. Switches use forwarding rules to pattern-match on the headers of

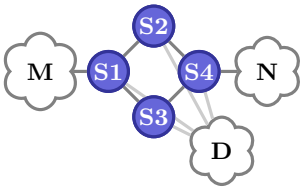

(a)

(b)

Initial configuration:

	Pattern	Action
S1	in=M	out=S2
S2	in=S1,type=UDP	out=D
S2	in=S1	out=S4
S3	in=S1	out=S4
S4	in=S2	out=N

Final configuration:

	Pattern	Action
S1	in=M	out=S3
S2	in=S1	out=S4
S3	in=S1,type=UDP	out=D
S3	in=S1	out=S4
S4	in=S2	out=N

(c) $\varphi \triangleq (\text{port} = M \wedge \text{type} = \text{UDP}) \Rightarrow F(\text{port} = \text{D}) \wedge (\text{port} = M \wedge \text{type} \neq \text{UDP}) \Rightarrow F(\text{port} = N)$

Fig. 1: Example: (a) topology, (b) configurations, and (c) property.

incoming packets and forward them out one of the switch's ports. To implement a network-wide update from one version to another, the operator must modify the rules on individual switches. It is not possible to implement simultaneous coordinated updates to multiple switches in general, as the network is a distributed system. Hence, the sequence of update commands must be carefully constructed to ensure that important properties are not violated.

To illustrate, consider the example shown in Figure 1. The topology includes networks M, N, D, and switches S_j (for j from 1 to 4). The property φ states that the network should forward non-UDP packets from M to N, and drop UDP packets (by forwarding them to D). The initial and final configurations satisfy this property: the former forwards traffic along the upper path M, S_1, S_2, S_4, N, filtering UDP packets at S_2, while the latter forwards traffic along the lower path, filtering at S_3. However, blindly modifying rules on each switch can cause a property violation. Updating S_1 first allows UDP traffic to flow to N along the lower path. A correct update begins with the sequence S_3, S_1, S_2. However, finding such a sequence is a significant challenge for the operator, since it requires reasoning about intermediate configurations and their effect on in-flight packets. In contrast, our system finds the correct sequence automatically.

3 Preliminaries

To start, we review basic definitions related to temporal logic, develop a network model, and define the network update problem formally.

Kripke structures. A *Kripke structure* is a tuple $(S, S_0, \delta, \lambda)$, where S is a finite set of states, $S_0 \subseteq S$ is a set of initial states, $\delta \subseteq S \times S$ is a transition relation, and the function $\lambda : S \to 2^{AP}$ labels each state with a set of atomic propositions drawn from a fixed set AP. A Kripke structure is *complete* if every state has at least one successor. A state $s \in S$ is a *sink state* if for all states s', $\delta(s, s')$ implies that $s = s'$, and *loop-free* if the only cycles are self-loops on sink states. In this paper, we consider complete and loop-free Kripke structures. A *trace* t is an infinite sequence of states, $t_0, t_1, \ldots$ such that $\forall i \geq 0 : \delta(t_i, t_{i+1})$. Given a trace t, t^i is the suffix of t starting at the i-th position—i.e., $t^i = t_i t_{i+1} \ldots$. Given a set of traces T, we let T^i denote the set $\{t^i \mid t \in T\}$.

Formulas. We define Linear Temporal Logic (LTL) formulas over a fixed set of atomic propositions AP. An LTL formula φ in negation normal form (NNF) is either an atomic proposition p in AP, negated atomic proposition $\neg p$, a disjunction $\varphi_1 \vee \varphi_2$, conjunction $\varphi_2 \wedge \varphi_2$, next $X\varphi$, until $\varphi_1 U \varphi_2$, or release $\varphi_1 R \varphi_2$, where φ_1 and φ_2 are LTL formulas in NNF. The semantics of LTL formulas is defined on traces of Kripke structures in a standard way. The operators F and G can be defined using the other connectives. We write $t \models \varphi$ to indicate that the trace t satisfies the formula φ and $T \models \varphi$ to indicate that $t \models \varphi$ for each t in T. Given a state s of a Kripke structure K, let $\mathit{traces}_K(s)$ be the set of traces of K starting from s and $\mathit{succ}_K(s)$ be the set of states defined by $s' \in \mathit{succ}_K(s)$ if and only if $\delta(s, s')$. We will omit the subscript K when it is clear form the context. A Kripke structure $K = (S, S_0, \delta, \lambda)$ satisfies an LTL formula φ if for all states $s_0 \in S_0$ we have that $\mathit{traces}(s_0) \models \varphi$.

Network model. A *packet* pt is a set of key-value pairs, representing header fields such as IP addresses etc. We model these using atomic propositions such as (*protocol*=ICMP) and use LTL formulas to describe the paths that a packet may take through the network. More formally, we let AP_H be the set of all such propositions, and refer to the set 2^{AP_H} as *traffic classes*. Although real-world switches can modify header fields, we are interested in reasoning about forwarding, so we will model packet contents as fixed values. Given a packet $pt \subseteq AP_H$ and a propositional formula ϕ, we define satisfaction in the usual way, i.e. $pt \models (field{=}val)$ if $(field{=}val) \in pt$, and $pt \models \phi_1 \wedge \phi_2$ if $pt \models \phi_1 \wedge pt \models \phi_2$, and similarly for $pt \models \phi_1 \vee \phi_2$.

A switch s forwards packets pt based on the values of their header fields and the ingress port they arrive on. We model the *configuration* of a switch, written sp, as a set of *rules*, each of the form (k, n_i, ϕ, D) with $k \in \mathbb{N}$ a priority, $n_i \in \mathbb{N}$ an ingress port, ϕ a pattern represented as a propositional formula over AP_H, and $D \subseteq \mathbb{N}$ a set of egress ports. When a packet pt enters s via port n, the rules are processed in order of priority, and if a match (k, n, ϕ, D) is found such that $pt \models \phi$, the packet pt is copied to each of the destination ports in $n_o \in D$ and rule processing terminates. We use $\mathit{forwards}(s, pt, n, n_o)$ to indicate that pt is forwarded by s from n to n_o in this way. We model a *network configuration* nc as a set of switch configurations, and we use $\mathit{forwards}(nc, pt, n, n_o)$ as shorthand for $\exists s \in S, \mathit{forwards}(s, pt, n, n_o)$.

A *topology* T is a tuple (S, P_i, P_o, f, L, I) where S is a set of switches, P_i, P_o are sets of ingress and egress ports respectively, $f : P_i \cup P_o \to S$ is a (bijective) mapping from ports to switches, $L \subseteq (P_o \times P_i) \cup \{(n, n) | n \in P_o\}$ is a set of physical links, and $I \subseteq P_i$ is a set of *traffic sources*, i.e. locations where packets can enter the network. There is a special port *drop* used to model switches dropping packets. Without loss of generality, we stipulate that egress ports must always have outgoing links—i.e. $\forall n_o \in P_o, \exists n \in P_i \cup P_o$ such that $(n_o, n) \in L$. A disconnected egress port n_o can be modeled with a link (n_o, n_o).

If there is a link $(n_{out}, n_{in}) \in L$, then a packet at n_{out} will move to n_{in}—i.e. nc and L induce a set $\mathit{traces}(L, nc)$ of *packet traces* describing how packets can travel through the network. A trace tr is a sequence $lp_0, lp_1, \cdots, lp_m, lp_m, \cdots$,

	Incoming,	Condition		Outgoing
1	$port = 1,$	$type = SSH$	$\longrightarrow$	$port' = 5, 4$
2	$port = 1,$	$src = XYZ \vee$		
		$type = SSH$	$\longrightarrow$	$port' = 4$
3	$port = 2,$	$true$	$\longrightarrow$	$port' = 3$

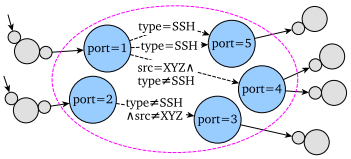

Fig. 2: Sample flow table and representation as a Finite-State Machine

where each lp_k is a tuple (pt, n_k, t_k) indicating that packet pt is located at port $n_k \in \mathbb{N}$ at time $t_k \in \mathbb{R}$. The timestamp will be used to simplify the discussion of temporal ordering of network updates. The initial lp_0 is located at a traffic source (i.e. $lp_0 = (pt, n, t)$ for some $n \in I$), the final repeating lp_m is located at a traffic sink, and $\forall k \in \{0, 1, \cdots, m-1\}, \big(t_{k+1} > t_k \wedge ((n_k, n_{k+1}) \in L \vee \mathit{forwards}(nc, tr, n_k, n_{k+1}))\big)$. Although there may be physical loops in the network topology, we require that all traces are *loop-free* in the same way as traces in Kripke structures.

A *switch update* u is a pair (s, sp) comprising a switch s and a switch configuration sp. Given a network configuration nc and a switch update (s, sp), the expression $nc[s \leftarrow sp]$ denotes the network configuration nc', where $nc'(s) = sp$ and $nc'(s') = nc(s')$ if $s' \neq s$.

In some situations, if we implement updates at the granularity of whole switch configurations, there is no correct update sequence. Our formal network model used in this paper only considers updates to whole switches—*switch granularity*. However, this is not a limitation: updates that operate at finer granularity, such as updating individual rules, can be easily reduced to switch granularity by adding additional switches to represent each rule. Our system assumes that this reduction has already been performed if needed.

We can transform a topology T and network configuration nc into a Finite-State Machine (FSM) F suitable for model checking. Each port becomes a state in F, each link becomes an unlabeled directed edge, and each rule (k, n_i, ϕ, D) becomes a set of directed edges $\{(n_i, n_o) | n_o \in D\}$, each labeled with the propositional formula $\phi \wedge (\bigwedge_{j=1}^{k-1} \neg\phi_j)$ obtained from the rule pattern and the negations of previous rule patterns (see Figure 2). We then convert this FSM into a Kripke structure M by constructing a Kripke structure M_k for each traffic class T_k (by taking only edges $e \in F$ where $T_k \Rightarrow e$) and letting $M = \bigcup M_k$.

A *command* c is a pair (u, t) where u is an update (s, sp), and $t \in \mathbb{R}$ is a timestamp. The time parameter allows commands to be spaced far enough apart to ensure that each packet experiences at most a single network update. Intuitively, waiting "long enough" makes sense only for networks where every packet arrives at a sink in bounded time; we call such networks *wait-correct*, and we call the bound *wait*. Real-world networks have this property either because their configurations are loop-free, or due to time-to-live fields. We call sequences where a delay of at least *wait* exists between every switch update a *careful*

sequence. We will consider only such sequences, and we present a sound and complete algorithm for efficiently finding them.

Let $CS = c_0, c_1, \ldots, c_{k-1}$ be a sequence of commands, and let nc be a configuration. A sequence of network configurations $nc_0, nc_1, \ldots, nc_n$ is induced by CS and nc, if $nc_0 = nc$, and for all j between 0 and $k-1$, if $c_j = ((s, sp), t)$ then $nc_{j+1} = nc_j[s \leftarrow sp]$. We write $nc \xrightarrow{CS} nc'$ if the last element of the sequence induced by CS is nc'—i.e. CS starts with nc and produces a sequence ending with nc'. We wish to define the set $\mathit{traces}(CS, nc)$ containing all possible traces a packet could take through the network, regardless of when the update commands are executed. Note that this is a superset of the traces induced by all nc_j. Let $t_0, t_1, \cdots, t_{k-1}$ be the times corresponding to the commands in CS. We write $\mathit{configs}(nc, t)$ to denote the set of configurations occurring immediately before t—i.e. $\{nc_j \mid j = max(\{j \mid 0 < j \leq k-1 \wedge t_j < t\})$, and for a set of times τ, we let $\mathit{configs}(nc, \tau)$ denote $\bigcup_{t \in \tau} \mathit{configs}(nc, t)$. For trace $tr = lp_0, lp_1, \cdots, lp_m, lp_m, \cdots$, we say that CS *allows* tr if both of the following hold:

- the sink state $lp_m{=}(pt, n_m, t)$ satisfies $(n_m, n_m) \in L$—i.e. there is a self-loop
- for non-sink states $lp_j{=}(pt, n_1, t_\alpha)$, if $lp_{j+1}{=}(pt, n_2, t_\beta)$, then $((n_1, n_2) \in L)$ or $(\exists nc_k \in \mathit{configs}(nc, [t_\alpha, t_\beta])$ such that $\mathit{forwards}(nc_k, pt, n_1, n_2))$.

We use $\mathit{traces}(CS, nc)$ to denote the set of allowed traces. This set contains all possible paths a packet can take through the network, regardless of when the update commands occur in relation to the (in-flight) packet.

A sequence is *simple* if each switch appears at most once in the updates. A command sequence CS is *correct* with respect to a formula φ and configuration nc_i if we have $\forall tr \in \mathit{traces}(CS, nc_i) : tr \models \varphi$. We will show that checking this correctness can be done by checking the correctness of each nc_j in the sequence.

We are now ready to state the *network update synthesis problem* formally. Given an initial network configuration nc_i, a final network configuration nc_f, and an LTL specification φ, construct a careful, simple sequence of commands CS such that (a.) $nc_i \xrightarrow{CS} nc_f$, and (b.) CS is correct with respect to φ and nc_i.

4 Update Synthesis Algorithm

This section presents an algorithm for synthesizing network updates. It uses depth-first search (DFS) to explore the space of possible solutions, uses counterexamples to quickly detect erroneous configurations, and incorporates heuristics to avoid backtracking when possible. We prove that our algorithm is sound and complete for the class of simple and careful updates.

Algorithm. Figure 3 presents the OrderUpdate algorithm in pseudocode. It either returns a simple sequence of updates or fails if no such sequence exists. Most of the actual work is done by DFSForOrder, which manages the search and invokes the model checker. We use DFS rather than another traversal strategy because we typically expect many update sequences to be valid.

Each call to DFSForOrder attempts to add one switch to the current update sequence, yielding a new network configuration. We maintain two formulas, V

Procedure OrderUpdate(nc_i, nc_f, φ)
Input: Initial configuration nc_i, final configuration nc_f, and formula φ.
Output: Simple and careful update sequence L, or an error message if no such update exists

```
1: W ← false
2: V ← false
3: (ok, L) ← DFSforOrder(nc_i, ⊥,φ)
4: if ok then
5:     return L
6: else
7:     return "No update exists."
```

Procedure DFSforOrder(nc, s,φ)
Input: Current configuration nc, next switch to update s, and formula φ.
Output: Boolean ok if a correct update exists; correct update sequence L

```
1: if nc ⊨ V ∨ W then return (false,[])
2: (ok,cex) ← modelCheck(nc,s,φ)
3: W ← W ∨ analyzeCex(cex)
4: V ← V ∨ nc
5: if ¬ ok then return (false,[])
6: if nc = nc_f then return (true,[nc])
7: for (nc', s') ∈ nextPolicies(nc) do
8:     (ok,L) ← DFSforOrder(nc',s',φ)
9:     if ok then return (true,nc :: L)
10: return (false,[])
```

Fig. 3: OrderUpdate Algorithm.

and W, that track the set of configurations that have been visited so far and the set of configurations excluded by counterexamples respectively. To check whether all packet traces traversing this configuration satisfy the LTL specification φ, we call ModelCheck(nc,φ,s) (line 2 of DFSforOrder). The model checker takes as input the switch s being updated. This is to facilitate performing model checking using incremental techniques (see Section 5). We also assume that the model checker checks whether the new configuration has loops.

If the model checker returns true, then the nc is safe and the search for a solution proceeds recursively. Otherwise, the resulting counterexample is analyzed, and added to W. The function analyzeCex(cex) analyzes the counterexample cex, and returns a formula representing the set of switches that occurred in the counterexample trace, and flags indicating whether each switch was updated or not. This allows equivalent future configurations to be immediately eliminated without invoking the model checker. To illustrate, recall Figure 1 and suppose that we first update switch S_2. The resulting configuration is incorrect, as there could be a non-UDP packet that reaches N via the trace S_1, S_2, S_4. The information that configurations where S_1 not updated, S_2 updated, and S_4 not updated is unsafe would then be added to W. In practice, many counterexamples consist of only a small number of nodes compared to the total size of the network.

Formal Properties. To show that OrderUpdate is sound, we prove that if the algorithm returns a sequence of updates, then this sequence is correct with respect to the specification φ and the initial configuration nc_i given as input. To save space, we defer the proof (and all other proofs) to Appedix C.

Theorem 1 (Soundness). *Given an initial configuration nc_i a final configuration nc_f, and an LTL formula φ,* OrderUpdate *returns a command sequence CS, then $nc_i \xrightarrow{CS} nc_f$, and CS is correct with respect to φ and nc_i.*

To show that OrderUpdate is complete with respect to simple command sequences, we observe that OrderUpdate searches through all simple sequences.

Theorem 2 (Completeness). *Given an initial configuration nc_i, a final configuration nc_f, and a specification φ, if there exists a simple sequence of commands CS with $nc_i \xrightarrow{CS} nc_f$, then* OrderUpdate *returns one such sequence.*

Optimizations. To avoid backtracking during DFS, we introduce a heuristic that leverages partial results computed by the model checking algorithm (see Section 5), which labels the underlying Kripke structure by LTL formulas. Our heuristic biases for updates that minimize the number of changes to the labeling. Intuitively, we prefer updates that cause "minimal waves"—i.e., updates whose re-labelings propagate the least distance—as such updates can be checked quickly, and hint at the fact that the property of interest is preserved. For example, recall Figure 1. The initial labeling labels S_1 and S_2 by type $\neq$ UDP $\Rightarrow F(\text{port} = \text{N})$. If we update S_2 first, we would need to update the states representing both S_2 and S_1 to propagate the formula $F(\text{port} = \text{N})$. Updating S_1 first would have a similar effect. However, if we update S_3 first, no other states are affected. After S_3 is updated, S_1 can be updated with minimal effect. Finally, S_2 can be updated. This shows how model checking information can be exploited to obtain a useful synthesis heuristic.

5 Incremental Model Checking

We now present an incremental algorithm for model checking Kripke structures. This algorithm is central to our synthesis tool, which invokes the model checker on many closely related structures as it computes updates. The algorithm assumes that the only cycles in the Kripke structure are self-loops on sink nodes—something that is true of structures that encode loop-free network configurations. It works by re-labeling the states of a previously labeled Kripke structure with the (possibly different) set of formulas that hold after a modification.

State Labeling. We begin by presenting a simple algorithm for labeling states of a Kripke structure with sets of formulas, following the approach introduced by Wolper, Vardi, and Sistla (WVS) [35]. The WVS algorithm translates an LTL formula φ into a local automaton and an eventuality automaton. The local automaton checks the consistency between a state and its parent, and handles labeling of all formulas except $\varphi_1 \; U \; \varphi_2$, which is checked by the eventuality automaton. The two automata are then composed into a single Büchi automaton A whose states correspond to subsets of the set of subformulas of φ and their negations. Hence, we can label each state of the Kripke structure by a set L of sets of formulas so that if a state s is labeled by L, then for each S in L, there exists a trace t starting from s satisfying all the formulas in S. When we define $ecl(\varphi)$, we will see that these are exactly the formulas in $ecl(\varphi)$ that t satisfies.

We now describe state labeling precisely, deferring the definitions of several auxiliary functions to Appendix A. Let φ be an LTL formula in NNF. The

extended closure of φ, written $ecl(\varphi)$, is the set of all subformulas of φ and their negations. A subset $M \subset ecl(\varphi)$ of the extended closure is said to be *maximally consistent* if it contains $\top$ and is simultaneously closed and consistent under boolean operations (e.g. $\varphi_1 \vee \varphi_2 \in M$ if and only if $\varphi_1 \in M$ or $\varphi_2 \in M$). Likewise, the function $follows(M_1, M_2)$ captures the notion of successor induced by LTL's temporal path operators, lifted to maximally-consistent sets (e.g., $X\ \varphi_1 \in M_1$ iff $\varphi_1 \in M_2$). Given a trace t and a maximally-consistent set M, we write $t \models M$ if and only if for all $\psi \in M$, we have $t \models \psi$.

For the rest of this section, we fix a Kripke structure $K = (S, S_0, \delta, \lambda)$, a state s in S, an LTL formula φ in NNF, and a maximally-consistent set $M \subset ecl(\varphi)$.

To compute the label of a state s, there are two cases depending on whether it is a sink state or a non-sink state. If s is a sink state, the function $HoldsSink(s, M)$ computes a predicate that is true if and only if, for all $\psi \in M$ and the unique trace t starting from s, we have $t \models \psi$. More formally, $HoldsSink(s, M)$ is defined to be $(\forall \psi \in M : Holds_0(s, \psi))$. The auxiliary function $Holds_0$, defined in the appendix, computes a predicate that is true if and only if ψ holds at s. For example, $Holds_0(s, \phi_1\ \mathrm{U}\ \phi_2)$ is defined as $Holds_0(s, \phi_2)$ because the only transition from s is a self-loop.

For the second case, suppose s is a non-sink state. If we are given a labeling for $succ_K(s)$, we can easily extend it to a labeling for s. Let $V \subseteq S$ be a set of vertices. A function $labelGraph_K$ is a *correct labeling* of K with respect to φ and V if for every $v \in V$, it returns a set L of maximally consistent sets such that $M \in L$ if and only if $M \subseteq ecl(\varphi)$ and there exists a trace t in $traces(v)$ such that $t \models M$. Suppose that $labelGraph_K$ is a correct labeling of K with respect to φ and $succ_K(s)$. The function $Holds_K(s, M, labelGraph_K)$ computes a predicate that is true if and only if there exists a trace t in $traces_K(s)$ with $t \models M$. More formally, the function $Holds_K(s, M, labelGraph_K)$ is defined as $(\lambda(s) = (AP \cap M)) \ \wedge\ \exists s' \in succ_K(s), M' \in labelGraph_K(s') : follows(M, M')$.

The following lemmas capture the correctness of these labeling operations:

Lemma 1. $HoldsSink(s, M) \iff \exists t \in traces(s) : t \models M$.

Lemma 2. *If* $labelGraph_K$ *is a correct labeling with respect to* φ *and* $succ_K(s)$*, then* $Holds_K(s, M, labelGraph_K) \iff \exists t \in traces_K(s) : t \models M$.

Finally, we define $labelNode_K(\varphi, s, labelGraph_K)$, which computes a label L for s such that $M \in L$ if and only if there exists a trace $t \in traces_K(s)$ such that $t \models M$ for all $M \subset ecl(\varphi)$. We assume that $labelGraph_K$ is a correct labeling of K with respect to φ and $succ(s)$. For sink states, $labelNode_K(\varphi, s, labelGraph_K)$ returns $\{M \mid M \in ecl(\varphi) \wedge HoldsSink(s, M)\}$, while for non-sink states it returns $\{M \mid M \in ecl(\varphi) \wedge Holds_K(s, M, labelGraph_K)\}$.

Incremental algorithm. To incrementally model check a modified Kripke structure, we must re-label its states with the formulas that hold post-update. Assume that the update u to the Kripke structure is given by a set of nodes $S'_0 \subseteq S$, a set of nodes $S' \subseteq S$, and a relation $\delta' \subseteq S' \times S$. We can apply u to K to obtain a modified Kripke structure $upd(K, u)$ as follows: $K' = (S, S'_0, \delta'', \lambda)$, where

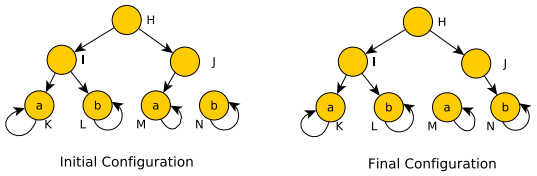

Fig. 4: Illustration of the incremental algorithm.

$(s_1, s_2) \in \delta''$ if and only if (i) $s_1 \in S'$ and $(s_1, s_2) \in \delta'$, or (ii) $s_1 \in (S \setminus S')$ and $(s_1, s_2) \in \delta$. An update might add or remove edges connected to a small set of nodes, corresponding to a change in the rules on a switch. Suppose that $labelGraph_K$ is a correct labeling of K with respect to φ and S. The incremental model checking problem is to determine whether $K' = upd(K, u)$ is correct with respect to φ. The naive approach is to model check K_1 without using the labeling $labelGraph_K$. We call this the *monolithic* approach. In contrast, the *incremental* approach re-uses the results of model checking K to efficiently verify K'.

Example. To illustrate, consider the structure on the left side of Figure 4, with H the only initial state. Suppose that the update modifies J, and the δ' relation only contains the pair (J, N). The resulting structure is shown on the right side of Figure 4. Consider labeling the structure with formulas $F\ a$, $F\ b$, and $F\ a \vee F\ b$. On the left hand side, we will have that the nodes K and M are labeled by $F\ a$, the nodes L and N by $F\ b$, the nodes I and H by $F\ a \vee F\ b$, and the node J by $F\ a$. Also assume that we want to label the Kripke structure after the update on the right-hand side. Given that the update changes only the node J, the labeling can only change for J and its ancestors. We therefore start labeling node J, and find that its label should become $F\ b$. The labeling proceeds to node H, and finds that its label has not changed, it is still $F\ a \vee F\ b$. The labeling process could then stop, even if node H had ancestors.

Re-labeling states. Let $ancestors_K(V)$ be the set of ancestors of V in K—i.e., a set of vertices such that $ancestors_K(V) \subseteq S$ and $s \in ancestors_K(V)$ some node $v \in V$ is reachable from s. To define the incremental model checking algorithm for φ, we will need a function that takes as input a specification φ, a set of vertices V, a labeling $labelGraph_K$ that is correct for K with respect to φ and $S \setminus ancestors(V)$, and returns a correct labeling of K with respect to φ and S:

$$
\begin{aligned}
&relabel_K(\varphi, labelGraph, V) = \\
&\quad \begin{cases} labelGraph & \text{if } V = \emptyset \\ relabel_K(\varphi, labelGraph', V') & \text{otherwise} \end{cases} \\
&\quad \text{where } labelGraph'\ v = \begin{cases} labelNode_K(\varphi, v, labelGraph) & \text{if } v \in V \\ labelGraph(v) & \text{otherwise} \end{cases} \\
&\qquad \text{and } V' = \{s \mid \exists v \in V : v \in succ_K(s)\}
\end{aligned}
$$

Theorem 3. *Let $V \subseteq S$ be a set of vertices and $labelGraph_K$ a correct labeling with respect to φ and $S \setminus ancestors_K(V)$. Then $relabel_K(\varphi, labelGraph_K, V)$ is a correct labeling with respect to φ and S.*

Given a labeling that is correct with respect to φ and S, it is easy to check whether φ is true for all the traces starting in the initial states: the predicate $checkInitStates_K(labelGraph_K, \varphi)$ is defined as $\forall s_0 \in S_0, M \in labelGraph(s) : \varphi \in M$. Next, let S_f be the set of all sink states of K. Then $ancestors_K(S_f)$ is the set S of all states K. Therefore, for any initial labeling lG_0, $relabel(\varphi, lG_0, S_f)$ is a correct labeling with respect to φ and S. The function $modelCheck_K(\varphi)$ is defined to be equal to $checkInitStates_K(relabel_K(\varphi, lG_0, S_f))$, where we can define lG_0 to be the empty labeling $\lambda v.\emptyset$.

Corollary 1. $modelCheck_K(\varphi) = true \iff K \models \varphi$.

We now define our incremental model checking function. Let $u = (S'_0, S', \delta')$ be an update, $K' = upd(K, u)$ the updated Kripke structure, and $labelGraph_K$ a previously-computed correct labeling of K with respect to φ and S, the states of K. The function $incrModelCheck(K, u, \varphi, labelGraph_K)$ is defined as $checkInitStates_{K'}(relabel_{K'}(\varphi, labelGraph_K, S'))$.

Corollary 2. $incrModelCheck(K, u, \varphi, labelGraph) = true \iff K' \models \varphi$

The runtime complexity of the *modelCheck* function is $O(|K| \times 2^{|\varphi|})$. The runtime complexity of *incrModelCheck* is $O(|ancestors_K(V)| \times 2^{|\varphi|})$, where V is the set of nodes being updated. Note that to achieve this complexity, the function *labelNode* must be implemented efficiently.

Counterexamples. The incremental algorithm just described can be extended to generate counterexamples in cases where the formula does not hold. A formula $\neg\varphi$ does not hold, if an initial state is labeled by L, such that there exists a set $M \in L$, such that $\neg\varphi \in M$. Examining the definition of *labelNode*, we find that in order to add a set M to the label L of a node s, there is a set M' in the label of one a child s' of s that explains why M is in L. The first node of the counterexample trace starting from s is one such child s'.

6 Implementation and Experiments

We built a full working implementation of our update synthesis tool in OCaml. It consists of 5KLoc that implements the synthesis algorithm described in Section 4, the incremental model checker described in Section 5, and several interfaces for specifying network updates including a data type (T, nc_i, nc_f, φ) for describing problem instances, a simple surface syntax, and a front-end that parses topology information, which can be used to build randomized update problems. The tool works by first building a Kripke structure using the algorithm described in Section 3 and then communicating with the model checker *MC* by modifying the

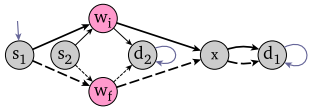

Fig. 5: Waypointing example.

graph to indicate which switches have updated so far. We provide three model checking backends: *Incremental* uses the incremental algorithm to check and recheck formulas, *Monolithic* uses a labeling technique but re-labels the entire graph on each call, and *NuSMV* encodes the FSM representation of the current network configuration as a transition relation and queries NuSMV. All three backends learn from counterexamples, as described in Section 4.

Experiments. To evaluate the performance of our tool, we conducted experiments on real-world and synthetic network topologies. Our experiments address two questions: (1) whether our tool can be used to solve update problems of realistic size, and (2) how it scales as inputs grow. We gathered real and synthetic network topologies, ranging in size from small (5-10 switches) to large (1000+ switches). We used the *Topology Zoo* [18] dataset, which features 261 actual network topologies from around the world, as well as synthetically constructed *Fat Tree* [21] and *Small World* topologies [25], which are often used in datacenters. We then built configurations for these topologies, and constructed formulas that hold in the initial/final configurations. We then ran our synthesizer and measured the amount of time needed to generate the update sequence.

To create configurations for each topology, we first computed all-pairs shortest path information for the switches $S \in T$ and then randomly partitioned the switches into sources and destinations, and selected paths P_i and P_f for use in the initial and final configurations. In addition, some paths were randomly selected to waypoint through nodes W_i and W_f respectively as shown in Figure 5.

As properties for the update, we used the following LTL formulas:

- *Reachability:* traffic from a given source must reach a certain destination:
$$(\text{port} = s1) \Rightarrow \text{F}\,(\text{port} = d1)$$
- *Waypointing:* traffic must traverse a waypoint w:
$$(\text{port} = s1) \Rightarrow ((\text{port} \neq d1)\,\text{U}\,((\text{port} = w_i) \wedge \text{F}\,(\mathit{port} = d1)))$$
- *Service chaining:* traffic must waypoint through several intermediate nodes:
$$(\text{port} = s1) \Rightarrow ((\text{port} \neq x \wedge \text{port} \neq d1)\,\text{U}\,((\text{port} = w_i) \wedge ((\text{port} \neq d1)\,\text{U}\,((\text{port} = x) \wedge \text{F}\,(\mathit{port} = d1)))))$$

Incremental vs. monolithic. Our first experiment compared the performance of Incremental against NuSMV. The results are shown in Figure 6. Of all 261 models in the Topology Zoo, our tool solved $238/261 \approx 91\%$ while NuSMV solved $212/261 \approx 81\%$ before the 5-minute timeout expired. Of the ones solved by both, the answers were consistent, and Incremental was faster $185/212 \approx 87\%$,

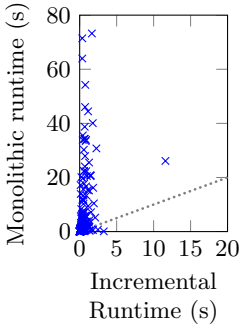

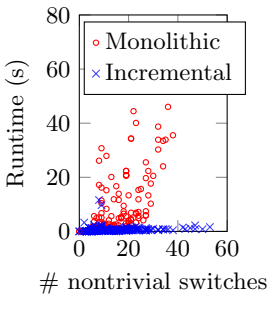

Total	Count	Incr.	NuSMV	Win
212	186	YES	YES	160
	26	NO	NO	25
12	4	YES	timeout	4
	8	NO	timeout	8
14	9	YES	crash	9
	5	NO	crash	5
21	21	timeout	timeout	0
2	2	timeout	crash	0

Summary of answers

Fig. 6: Topology Zoo Performance of Incremental vs. NuSMV

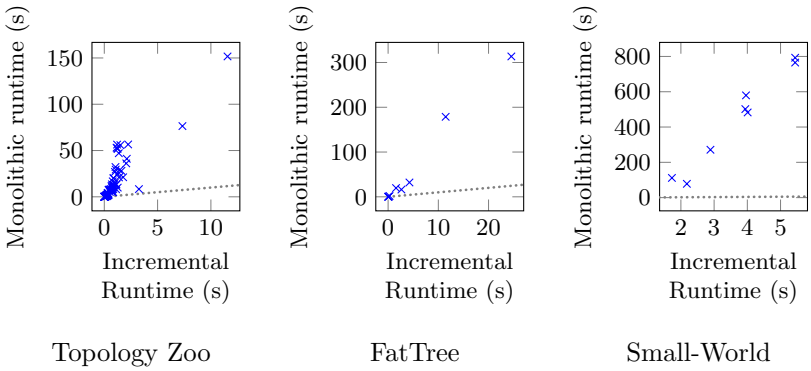

Fig. 7: Incremental vs. Monolithic

with speedups up to 203.5x, and 17.1x on average. This experiment clearly demonstrates the advantages of using the Incremental algorithm.

We also compared the performance of the Incremental and Monolithic backends. Monolithic relabels the entire graph during each model-checking call, incurring a large performance penalty. Incremental performed better on all examples, with average speedup of 13.2x, 7.9x, 109.3x on the three datasets shown in Figure 7. The maximum runtimes using Incremental were 11.6s, 147.5s, and 21.0s respectively. The grey line depicts $y = x$—i.e., equivalent performance—but has a low slope due to the axis scales. Monolithic is almost always above the line.

Scalability To quantify the scalability of our system, we constructed Small World topologies up to 1500 switches, and conducted experiments with all three types of properties, obtaining the results shown in Figure 8. The largest example had 1131 switches with nontrivial updates . The maximum graph generation time was 116.5s, and the maximum synthesis time was 28.0s. This experiment shows that our implementation is able to scale to networks of realistic size.

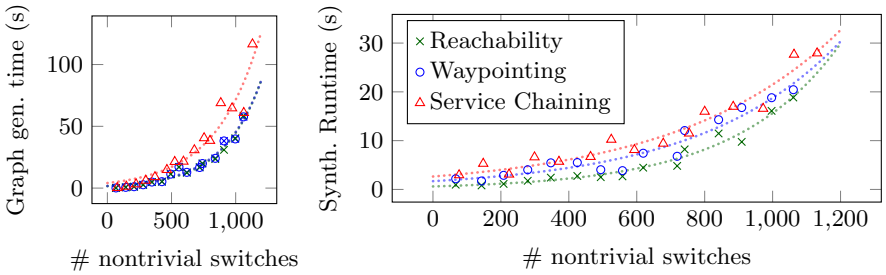

Fig. 8: Small-World Scalability of Incremental Algorithm

Optimizations. As described in Section 5, the incremental nature of our synthesis algorithm admits heuristics for guiding the DFS. This heuristic did not always help the overall *run-time* of our tool, but it did result in a reduction in the total number of model-checking calls for many examples with nontrivial updates. We randomly selected a subset of 84 Topology Zoo models, and found that 22 required nontrivial updates. The heuristic reduced the number of model-checker calls on 13, of which 12 were nontrivial ($12/22 \approx 55\%$), and maintained the same number of calls on the others. A larger performance gain was obtained on the Fat Tree examples, where 8 of the 9 examples were nontrivial. Out of these, the heuristic reduced total number of calls on 6 ($6/8 \approx 75\%$), and maintained the same number on the others. For the largest example (1023 switches with 511 nontrivial), the synthesizer runtime was reduced from 24.5s to 14.2s. We believe this heuristic could lead to better performance with further optimization.

7 Related Work

This paper extends preliminary work on using an LTL model checker to synthesize network updates, which was reported in a workshop paper [26]. Compared to that tool, our system replaces calls to an expensive external model checker with calls to a new built-in *incremental* model checker that we built from scratch. We have also extended the simple DFS search synthesis procedure with heuristics which allow it to find a solution quicker. Finally, we have evaluated the performance our tool on a comprehensive set of benchmarks with real-world topologies and policies. The rest of this section briefly reviews relevant work in network updates and model checking.

Synthesis of concurrent programs There is a large body of previous work on using synthesis to development of concurrent programs [5,34,31,33,33,13]. Work by Solar-Lezama et al. [31] and Vechev et al. [34], synthesizes a sequence of instructions, which is a similar task to ours. But traditional synthesis work and networking are quite different. First, traditional synthesis is a game against the environment which (in the concurrent programming case) provides inputs and schedules threads; in contrast our synthesis problem is a reachability problem on

the space of configurations. Second, the space of network configurations is very rich; checking properties is an LTL model checking problem in itself.

Network updates. The networking community has developed many protocol- and property-specific algorithms for implementing network updates. There are a number of systems for avoiding packet and bandwidth loss during planned maintenance to BGP [12,11,27,20]. Other work by Vanbever et al. avoids routing loops and blackholes during IGP migration [32]. In the context of software-defined networks, there are a number of projects that provide consistent network updates. Reitblatt et al. first proposed the notion of per-packet and per-flow consistent updates, as well as several implementation mechanisms, including two-phase updates [29]. Work by Katta et al. uses a partitioning scheme for propagating updates incrementally, which reduces the space overhead on switches. The SWAN system adds support for bandwidth guarantees [14]. A recent survey taxonomizes classes of consistent network updates [23]. Our work is different because we synthesize updates automatically from formal specifications.

Model checking. Model checking has been used to verify network properties in several recent systems [1,24,16,17]. Our network model resembles the internal representations used in these systems. The Incremental model checking problem has been studied previously [30,7,6]. Sokolsky and Smolka [30] present the first general incremental model checking algorithm for the alternation-free fragment of the modal μ-calculus. Their technique is based on computing the immediate effects of the update to the model, and then re-starting the fixpoint computation. We consider LTL specifications and specialize our algorithm to exploit the no-forwarding-loops assumption. Cohen and Kupferman [7] later extended the approach of Sokolsky and Smolka beyond the alternation-free fragment [19,3]. Chockler et al. [6] introduced an incremental model checking algorithm based on IC3 [4]. The algorithm re-uses both positive (proofs) and negative (counterexamples) results to speed up the model checking of the updated models. The algorithm is specific to the type of partial results produced by the IC3 algorithm. More generally, incremental model checking is an instance of *dynamic graph algorithms* [8]. The complexity bounds for such algorithms are open in many cases (for instance, for dynamic graph connectivity).

8 Conclusion

This paper describes a practical tool for automatically synthesizing correct network update sequences from formal specifications. Our approach is based on the observation that the search algorithm asks a series of related model checking questions. Exploiting the fact that correct configurations are loop-free, we build an efficient incremental model checker that performs orders of magnitude better than state-of-the-art monolithic tools. Experiments on real-world topologies demonstrate the effectiveness of our approach. In the future, we plan to explore extensions of our technique to deal with failures [22,28], and to use language-based representations to explore the foundations of network updates [2,9].

References

1. Ehab Al-Shaer and Saeed Al-Haj. FlowChecker: Configuration analysis and verification of federated OpenFlow infrastructures. In *SafeConfig*, 2010.
2. Carolyn Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger, and David Walker. NetKAT: Semantic foundations for networks. In *POPL*, January 2014.
3. R. Bloem, K. Ravi, and F. Somenzi. Efficient decision procedures for model checking of linear time logic properties. In *CAV*, 1999.
4. A. Bradley. SAT-based model checking without unrolling. In *VMCAI*, 2011.
5. Sigmund Cherem, Trishul Chilimbi, and Sumit Gulwani. Inferring locks for atomic sections. In *PLDI*, 2008.
6. H. Chockler, A. Ivrii, A. Matsliah, S. Moran, and Z. Nevo. Incremental formal verification of hardware. In *FMCAD*, pages 135–143, 2011.
7. G. Cohen and O. Kupferman. Incremental LTL model checking. In *Workshop on semantics and verification of hardware and software systems*, 2003.
8. D. Eppstein, Z. Galil, and G. Italiano. Dynamic graph algorithms. In *CRC Handbook*, 1997.
9. Nate Foster, Arjun Guha, Mark Reitblatt, Alec Story, Michael J Freedman, Naga Praveen Katta, Christopher Monsanto, Joshua Reich, Jennifer Rexford, Cole Schlesinger, et al. Languages for software-defined networks. *Communications Magazine, IEEE*, 51(2):128–134, 2013.
10. Pierre Francois and Olivier Bonaventure. Avoiding transient loops during the convergence of link-state routing protocols. *Networking, IEEE/ACM Transactions on*, 15(6):1280–1292, 2007.
11. Pierre Francois, Olivier Bonaventure, Bruno Decraene, and P-A Coste. Avoiding disruptions during maintenance operations on BGP sessions. *Network and Service Management, IEEE Transactions on*, 4(3):1–11, 2007.
12. Pierre Francois, Mike Shand, and Olivier Bonaventure. Disruption free topology reconfiguration in OSPF networks. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 89–97. IEEE, 2007.
13. Peter Hawkins, Alex Aiken, Kathleen Fisher, Martin Rinard, and Mooly Sagiv. Concurrent data representation synthesis. In *PLDI*, pages 417–428, June 2012.
14. Chi-Yao Hong, Srikanth Kandula, Ratul Mahajan, Ming Zhang, Vijay Gill, Mohan Nanduri, and Roger Wattenhofer. In *SIGCOMM*, pages 15–26, August 2012.
15. John P John, Ethan Katz-Bassett, Arvind Krishnamurthy, Thomas Anderson, and Arun Venkataramani. Consensus routing: The Internet as a distributed system. In *NSDI08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pages 351–364, 2008.
16. Peyman Kazemian, George Varghese, and Nick McKeown. Header space analysis: Static checking for networks. In *NSDI*, 2012.
17. Ahmed Khurshid, Wenxuan Zhou, Matthew Caesar, and P Godfrey. VeriFlow: Verifying network-wide invariants in real time. *ACM SIGCOMM Computer Communication Review*, pages 467–472, 2012.
18. S. Knight, H.X. Nguyen, N. Falkner, R. Bowden, and M. Roughan. *Selected Areas in Communications, IEEE Journal on, title=The Internet Topology Zoo*, 29(9):1765 –1775, october 2011.
19. O. Kupferman and M. Vardi. Relating linear and branching model checking. In *PROCOMET*, pages 304–326, 1998.

20. Nate Kushman, Srikanth Kandula, Dina Katabi, and Bruce M Maggs. R-BGP: Staying connected in a connected world. USENIX, 2007.
21. Charles E Leiserson. Fat-trees: universal networks for hardware-efficient supercomputing. *Computers, IEEE Transactions on*, 100(10):892–901, 1985.
22. Hongqiang Harry Liu, Xin Wu, Ming Zhang, Lihua Yuan, Roger Wattenhofer, and David Maltz. zUpdate: updating data center networks with zero loss. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, pages 411–422. ACM, 2013.
23. Ratul Mahajan and Roger Wattenhofer. On Consistent Updates in Software Defined Networks. In *SIGCOMM*, November 2013.
24. Haohui Mai, Ahmed Khurshid, Rachit Agarwal, Matthew Caesar, P Godfrey, and Samuel Talmadge King. Debugging the data plane with anteater. In *SIGCOMM*, pages 290–301, August 2011.
25. Mark EJ Newman, Steven H Strogatz, and Duncan J Watts. Random graphs with arbitrary degree distributions and their applications. *Physical Review E*, 64(2):026118, 2001.
26. Andrew Noyes, Todd Warszawski, and Nate Foster. Toward Synthesis of Network Updates. In *SYNT*, 2013.
27. Saqib Raza, Yuanbo Zhu, and Chen-Nee Chuah. Graceful network state migrations. *IEEE/ACM Transactions on Networking (TON)*, 19(4):1097–1110, 2011.
28. Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. FatTire: Declarative Fault Tolerance for Software-Defined Networks. In *HotSDN*, August 2013.
29. Mark Reitblatt, Nate Foster, Jennifer Rexford, Cole Schlesinger, and David Walker. Abstractions for network update. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 323–334. ACM, 2012.
30. O. Sokolsky and S. Smolka. Incremental model checking in the modal mu-calculus. In *CAV*, pages 351–363, 1994.
31. Armando Solar-Lezama, Christopher Grant Jones, and Rastislav Bodik. Sketching concurrent data structures. In *PLDI*, volume 43, pages 136–148, 2008.
32. Laurent Vanbever, Stefano Vissicchio, Cristel Pelsser, Pierre Francois, and Olivier Bonaventure. Seamless network-wide IGP migrations. In *SIGCOMM*, pages 314–325, 2011.
33. Martin Vechev and Eran Yahav. Deriving linearizable fine-grained concurrent objects. *ACM SIGPLAN Notices*, 43(6):125–135, 2008.
34. Martin Vechev, Eran Yahav, and Greta Yorsh. Abstraction-guided synthesis of synchronization. In *POPL*, volume 45, pages 327–338, 2010.
35. Pierre Wolper, Moshe Y. Vardi, and A. Prasad Sistla. Reasoning about infinite computation paths (extended abstract). In *FOCS*, pages 185–194, 1983.

A State Labeling Definitions

This appendix gives a number of auxiliary definitions related to state labeling.

Extended closure. The extended closure of a formula φ in NNF is defined as follows:
- $\top \in \mathit{ecl}(\varphi)$
- $\varphi \in \mathit{ecl}(\varphi)$
- If $\psi \in \mathit{ecl}(\varphi)$, then $\neg\psi \in \mathit{ecl}(\varphi)$. We identify ψ with $\neg\neg\psi$, for all ψ.
- If $\varphi_1 \vee \varphi_2 \in \mathit{ecl}(\varphi)$, then $\varphi_1 \in \mathit{ecl}(\varphi)$ and $\varphi_2 \in \mathit{ecl}(\varphi)$.
- If $\varphi_1 \wedge \varphi_2 \in \mathit{ecl}(\varphi)$, then $\varphi_1 \in \mathit{ecl}(\varphi)$ and $\varphi_2 \in \mathit{ecl}(\varphi)$.
- If $X\,\varphi_1 \in \mathit{ecl}(\varphi)$, then $\varphi_1 \in \mathit{ecl}(\varphi)$.
- If $\varphi_1\,U\,\varphi_2 \in \mathit{ecl}(\varphi)$, then $\varphi_1 \in \mathit{ecl}(\varphi)$ and $\varphi_2 \in \mathit{ecl}(\varphi)$
- If $\varphi_1\,R\,\varphi_2 \in \mathit{ecl}(\varphi)$, then $\varphi_1 \in \mathit{ecl}(\varphi)$ and $\varphi_2 \in \mathit{ecl}(\varphi)$.

Maximally-consistent set. A set $M \subseteq \mathit{ecl}(\phi)$ is said to be maximally-consistent if and only if it satisfies the following conditions:
- $\top \in M$
- $\psi \in M$ iff $\neg\psi \notin M$ (recall that we identify ψ with $\neg\neg\psi$, for all ψ)
- $\varphi_1 \vee \varphi_2 \in M$ iff ($\varphi_1 \in M$ or $\varphi_2 \in M$)
- $\varphi_1 \wedge \varphi_2 \in M$ iff ($\varphi_1 \in M$ and $\varphi_2 \in M$)

Follow function. The function $\mathit{follows}(M_1, M_2)$ is true if and only if M_1 and M_2 satisfy the following conditions:
- $X\ \varphi_1 \in M_1$ iff $\varphi_1 \in M_2$
- $\varphi_1\,U\,\varphi_2 \in M_1$ iff $\big(\varphi_2 \in M_1$ or $(\varphi_1 \in M_1$ and $\varphi_1\,U\,\varphi_2 \in M_2)\big)$
- $\varphi_1\,R\,\varphi_2 \in M_1$ iff $\big(\varphi_1 \in M_1$ or $(\varphi_2 \in M_1$ and $\varphi_1\,R\,\varphi_2 \in M_2)\big)$

Sink states. The function $\mathit{Holds}_0(s, \varphi)$ returns true if and only if for all traces t starting from s we have $t \models \phi$:

$$
\begin{aligned}
\mathit{Holds}_0(s, p) &= s \models p \\
\mathit{Holds}_0(s, \neg p) &= s \not\models p \\
\mathit{Holds}_0(s, \phi_1 \wedge \phi_2) &= \mathit{Holds}_0(s, \phi_1) \wedge \mathit{Holds}_0(s, \phi_2) \\
\mathit{Holds}_{0(}s, \phi_1 \vee \phi_2) &= \mathit{Holds}_0(s, \phi_1) \vee \mathit{Holds}_0(s, \phi_2) \\
\mathit{Holds}_0(s, \mathrm{X}\phi\ ) &= \mathit{Holds}_0(s, \phi) \\
\mathit{Holds}_0(s, \phi_1\ \mathrm{U}\ \phi_2) &= \mathit{Holds}_0(s, \phi_2) \\
\mathit{Holds}_0(s, \phi_1\ \mathrm{R}\ \phi_2) &= \mathit{Holds}_0(s, \phi_1) \vee \mathit{Holds}_0(s, \phi_2)
\end{aligned}
$$

B Appendix: Proof of correctness of incremental model checking.

Proof (Proposition 1). Observe that there is a unique trace t in $\mathit{traces}(s)$, as s is a sink state. We first prove that $t \models \varphi$ iff $\mathit{Holds}_0(s, \varphi)$. We prove this by induction on the structure of the LTL formula. Then we observe that there is a

unique maximally-consistent set M such that $t \models M$. This is the set $\{\psi \mid t \models \psi \wedge \psi \in \mathit{ecl}(\varphi)\}$. We then use the definition of $\mathit{HoldsSink}(s, M)$ for sink states to conclude the proof.

Proof (Theorem 2). We first prove soundness, i.e., if $\mathit{Holds}_K(s, M, \mathit{labelGraph}_K)$, then there exists $t \in \mathit{traces}(s)$ such that $t \models M$. We have $\mathit{Holds}_K(s, M, \mathit{labelGraph}_K)$ iff $(\lambda(s) = (AP \cap M))$ and there exists $s' \in \mathit{succ}_K(M)$, and $M' \in \mathit{labelGraph}_K(s')$ such that $\mathit{follows}(M, M')$. By assumption of the theorem, we have that if $M' \in \mathit{labelGraph}_K(s')$, then there exists a trace t' in $\mathit{traces}(s')$ such that $t' \models M'$. Consider a trace t such that $t_0 = s$ and $t^1 = t'$. For each formula $\psi \in M$, we can prove that $t \models \psi$, which finishes the soundness proof. The base case of the proof by induction is implied by the fact that $s \models (AP \cap M)$. The inductive cases are proven using the definitions of maximally-consistent set and the function *follows*.

We now prove completeness, i.e., that if there exists a trace t in $\mathit{traces}_K(s)$ such that $t \models M$, then $\mathit{Holds}_K(s, M, \mathit{labelGraph}_K)$ is true. Let t be the trace $ss_1s_2\ldots$. It is easy to see that if M is a maximally-consistent set, and $t \models M$, then $M = \{\psi \mid \psi \in \mathit{ecl}(\varphi) \wedge t \models \psi\}$. Let us consider the set of formulas $S = \{\psi \mid \psi \in \mathit{ecl}(\varphi) \wedge t^1 \models \psi\}$. Observe that S is a maximally-consistent set. By assumption of the theorem, we have that S is in $\mathit{labelGraph}_K(s_1)$. It is easy to verify that $\mathit{follows}(M, S)$. This concludes the completeness proof.

Proof (Theorem 3). We first note that only ancestors of nodes in V are re-labeled – all the other nodes are correctly labeled by assumption on *labelGraph*. We say that a node s is at level k with respect to a set of vertices T iff the longest simple path from s to a node in T is k. Let H_k be the set of nodes at level k from V. We prove by induction on k that at k-th iteration, $(\mathit{Sancestors}_K(V)) \cup H_k$ is a correct labeling of K w.r.t. φ and V. We can prove the inductive claim using Theorem 2 (for non-sink nodes) and Proposition 1 (for sink nodes).

Proof (Corollary 1). The result is obtained as a corollary of Theorem 3. Using this theorem, and the fact that the set $\mathit{ancestors}_K(S_f)$ is the set S of all states K, we obtain that $lG = \mathit{relabel}_K(\varphi, lG_0, S_f)$ is a correct labeling of K with respect to φ and S. In particular, for all initial states s_0, we have that for all $M \subset \mathit{ecl}(\varphi)$, $m \in lG(s_0)$ iff there exists a trace $t \in \mathit{traces}_K(s_0)$ such that $t \models M$. We now use the definition of *checkInitStates* to show that if *checkInitStates* returns true, then there is no initial state s_0 such that there exists $M \in lG(s_0)$ such that $\neg\varphi \in M$. Thus for all initial states s_0, for all traces t in $\mathit{traces}(t_0)$, we have that $t \models \varphi$.

Proof (Corollary 2).

Similar to the proof of Corollary 1.

C Appendix: Proofs of correctness for the synthesis algorithm

Soundness. Now we prove that OrderUpdate is sound. For the remainder of this section, let us fix a specific network topology (S, P_i, P_o, f, L, I).

A network configuration *nc induces* a one-packet trace t, if $tr \in traces(L, nc)$. A network configuration *nc* satisfies an LTL formula φ (denoted by $nc \models \varphi$), if for all packet traces *tr* induced by *nc*, we have that $tr \models \varphi$.

Let $CS = c_0 c_1 \dots c_{n-1}$ be a command sequence, and let *nc* be a network configuration. Let $nc_0 nc_1 \dots nc_n$ be the sequence of network configurations induced by *CS* and *nc*. Assume that *CS* is *careful* and *correct* with respect to an LTL formula φ and a network configuration *nc*.

Let nc_i be a network configuration. Let $tr = lp_0, lp_1, \dots lp_n, lp_n, \cdots$ be a packet trace contained in $traces(L, nc_i)$. Let $\sigma = s_0, s_1, \dots, s_{n-1}$ be the sequence of switches traversed by the packet, i.e. a sequence such that for $i < n$, we have that if $lp_i = (n, pt)$ and $lp_{i+1} = (n', pt')$, then $n \in f^{-1}(s_i) \cap P_i$ and $n' \in f^{-1}(s_i) \cap P_o$. Let $t_0, t_1, \cdots, t_n$ be the packet times. We fix *CS*, nc_i, *tr*, and σ for the following lemmas.

Our first lemma states that the path of every packet is affected by at most one update.

Lemma 3. *If CS is a careful update sequence, then* $N = configs(nc_i, [t_0, t_n])$ *contains at most 2 elements. If* $nc_j \in N$ *and* $nc_k \in N$*, then* $|j - k| \leq 1$.

Proof. Assume to the contrary that there are more than 2 elements. That is, there exist $t_0 \leq t_\alpha < t_\beta < t_\gamma \leq t_n$. Since *CS* is careful, we must have $t_\beta - t_\alpha \geq wait$ and $t_\gamma - t_\beta \geq wait$, meaning $t_\gamma - t_\alpha \geq 2 \cdot wait > wait$. But this means that $t_n - t_0 > wait$, i.e. the delay between lp_0 and lp_n is greater than *wait*, contradicting the definition of *wait*.

Let $nc_j \in N$ and $nc_k \in N$, and assume without loss of generality that $j \leq k$. If $nc_j = nc_k$, then clearly $j = k$ (i.e. $|j - k| = 0$). If $nc_j \neq nc_k$, then assume there is a w such that $j \leq w \leq k$. By definition of N, this forces $nc_w \in N$, and from the first part of the proof, we conclude that $nc_w = nc_j$ or $nc_w = nc_k$, meaning $w = j \vee w = k$. Thus, we must have $j + 1 = k$, i.e. $|j - k| = 1$.

The second lemma states that a path of every packet in the network could have occurred in one of the intermediate configurations; that is, no packet takes a path non-existent in any of the configurations, even though the packet might be in-flight during the updates.

Lemma 4. *If CS is careful and* $tr \in traces(CS, nc)$*, i.e. tr is allowed by CS, then* $tr \in traces(L, nc_j)$ *for one of the* nc_j *induced by CS and nc.*

Proof. From Lemma 3, we see that $N = configs(nc, [t_0, t_n])$ has at most 2 elements. If N is empty, then $t_0 = t_n$, since $t_n > t_0$ would require that there is at least one configuration before t_n. Then the definition of $traces(CS, nc)$ tells us that *tr* is simply a self-loop, which is a link in L, and hence $tr \in traces(L, nc)$. If there is a single element $nc \in N$, then the definition of $traces(CS, nc)$ tells us that each *lp* in *tr* traverses a link in L, or is forwarded by *nc*. This means $tr \in traces(L, nc)$.

If there are 2 elements, Lemma 3 tells us they are successive, i.e. nc_j, nc_{j+1}. This means that if $\sigma = s_0, s_1, \dots, s_{n-1}$ is the sequence of switches traversed by

tr, there is exactly one switch s_k whose configuration sp differs in nc_j, nc_{j+1}. Consider two adjacent terms in the trace, $lp_j = (pt, n_1, t_\alpha), lp_{j+1} = (pt, n_2, t_\beta)$. Since $tr \in traces(CS, nc)$, we have $(n_1, n_2) \in L$ or $forwards(nc_j, pt, n_1, n_2)$ or $forwards(nc_{j+1}, pt, n_1, n_2)$ by definition of allowed traces. But

$$forwards(nc_j, pt, n_1, n_2) \iff forwards(nc_{j+1}, pt, n_1, n_2)$$

unless $n_1, n_2 \in f^{-1}(s_k)$, since s_k is the only one whose configuration changed.

Thus, we have shown that for each pair of adjacent terms lp_j, lp_{j+1} in trace tr, either the packet travels along a link, or it is forwarded by *both* nc_j and nc_{j+1}, except for when the packet passes through s_k, in which case it forwarded by *either* nc_j and nc_{j+1}. From this, we conclude that $tr \in traces(L, nc_j)$ or $tr \in traces(L, nc_{j+1})$.

The third lemma states that carefulness (which is easily checkable) implies correctness.

Lemma 5. *If CS is careful and $nc_j \models \varphi$ for all the nc_j in the sequence induced by CS and nc_i, then CS is correct with respect to nc_i and φ.*

Proof. Let $tr \in traces(CS, nc_i)$. We can apply Lemma 3 to get $tr \in traces(L, nc_j)$ for some nc_j induced by CS and nc. Then by our hypothesis, $nc_j \models \varphi$, meaning $tr \models \varphi$, and since this is for arbitrary tr, CS is correct with respect to nc_i and φ.

Theorem 4 (Soundness). *If given an initial configuration nc_i a final configuration nc_f, and an LTL formula φ, OrderUpdate returns a command sequence CS, then*
- *$nc_i \xrightarrow{CS} nc_f$, and*
- *CS is correct with respect to φ and nc_i.*

Proof. It is easy to show that if OrderUpdate returns CS, then $nc_i \xrightarrow{CS} nc_f$. Each update in the returned sequence changes a switch configuration of one switch s to the configuration $nc_f(s)$, and the algorithm terminates when all (and only) switches s such that $nc_i(s) \neq nc_f(s)$ have been updated.

Now observe that if OrderUpdate returns CS, the sequence can be made careful by choosing an adequate time delay between each update command, and for all $j \in [0, n]$, $nc_j \models \varphi$. This is ensured by the call to a model checker (Line 2). We can use Lemma 5 to conclude that CS is correct with respect to φ and nc_i.

Completeness. The following proposition states that OrderUpdate is complete with respect to simple and careful command sequences. Let $CS = c_0, c_1, \ldots, c_{n-1}$ be a command sequence. Such a sequence is *simple* if for each $s \in \mathcal{S}$ there exists at most one i in $[0, n]$ such that c_i is an update of the form (s, sp). The proof (omitted here) uses the fact that OrderUpdate searches through all such sequences. The proposition characterizes the cases where the algorithm returns a solution.

Proposition 1. *Given an initial network configuration nc_i, a final network configuration nc_f, and a specification φ, if there exists a simple and careful sequence of commands CS such that $nc_i \xrightarrow{CS} nc_f$, then OrderUpdate returns one such sequence.*